

What is claimed is:

1 1. A method of efficiently storing content in a computing network, comprising steps of:
2 receiving hints regarding relationships among files; and
3 using the received hints to allocate storage for the files.

1 2. The method according to Claim 1, wherein the hints are created by a content management
2 system.

1 3. The method according to Claim 1, wherein the hints specify one or more files that are
2 likely to be referenced within a temporal proximity of a reference to a selected one of the files.

1 4. The method according to Claim 3, wherein the selected file is a web page.

1 5. The method according to Claim 4, wherein the one or more files comprise at least one of
2 (1) one or more embedded objects of the web page and (2) one or more other web pages which
3 are hyperlinked to the web page.

1 6. The method according to Claim 1, wherein the hints are created by a content authoring
2 tool, and wherein the hints specify one or more files that are likely to be referenced within a
3 temporal proximity of a reference to a selected one of the files.

1 7. The method according to Claim 6, wherein the selected file is a text document.

1 8. The method according to Claim 7, wherein the one or more files comprise one or more
2 objects which are embedded within or referenced by the text document.

1 9. The method according to Claim 3, wherein the hints further specify weights which
2 describe a degree of dependency for the relationships.

1 10. The method according to Claim 1, wherein the receiving step is performed by a file
2 system and the using step is performed by a storage system.

1 11. The method according to Claim 2, wherein the hints are encoded in a markup language
2 notation.

1 12. The method according to Claim 11, wherein the markup language notation is Extensible
2 Markup Language ("XML") notation.

1 13. The method according to Claim 1, further comprising steps of:
2 receiving a request for one of the files;
3 retrieving the requested file from the allocated storage; and
4 returning the retrieved file.

1 14. The method according to Claim 1, further comprising steps of:

2 using the received hints to create dependency information which is stored by a receiver of
3 the hints in temporary or permanent storage;
4 receiving a request for one of the files; and
5 determining a read request strategy for the requested file by accessing the stored
6 dependency information.

1 15. The method according to Claim 14, wherein the read request strategy comprises
2 determining selected ones of the files which should be pre-fetched along with a read of the
3 requested file.

1 16. The method according to Claim 15, wherein the step of determining selected ones further
2 comprises comparing a dependency weight of the files to a pre-fetch threshold.

1 17. The method according to Claim 16, wherein the pre-fetch threshold is used to tune the
2 pre-fetch operation.

1 18. The method according to Claim 15, further comprising steps of:
2 retrieving the requested file from the allocated storage;
3 retrieving the selected ones from the allocated storage;
4 returning the retrieved requested file; and
5 caching the retrieved selected ones.

1 19. The method according to Claim 18, further comprising the step of caching the retrieved
2 requested file.

1 20. A system for efficiently storing files in a computing network, comprising
2 means for receiving hints regarding relationships among files, wherein the hints specify
3 one or more files that are likely to be referenced within a temporal proximity of a reference to a
4 selected one of the files; and

5 means for using the received hints to allocate storage for the files.

1 21. A computer program product for efficiently storing files in a computing network, the
2 computer program product embodied on one or more computer-readable media and comprising:

3 computer readable program code means for receiving hints regarding relationships among
4 files, wherein the hints specify one or more files that are likely to be referenced within a temporal
5 proximity of a reference to a selected one of the files; and

6 computer readable program code means for using the received hints to allocate storage for
7 the files.